

Jet clustering with a scale-invariant filtered tree: A new algorithm for the substructure era

Andrew J. Larkoski¹, Denis Rathjens², Jason Veatch³, and Joel W. Walker⁴

¹*Department of Physics and Astronomy, University of California, Los Angeles, California 90095, USA*

²*Mitchell Institute for Fundamental Physics and Astronomy, Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843, USA*

³*Department of Physics, California State University East Bay, Hayward, California 94542, USA*

⁴*Department of Physics and Astronomy, Sam Houston State University, Huntsville, Texas 77341, USA*



(Received 10 April 2023; accepted 17 June 2023; published 11 July 2023)

We introduce a new jet clustering algorithm named SIFT (scale-invariant filtered tree) that maintains the resolution of substructure for collimated decay products at large boosts. The scale-invariant measure combines properties of k_T and anti- k_T by preferring the early association of soft radiation with a resilient hard axis, while avoiding the specification of a fixed cone size. Integrated filtering and variable-radius isolation criteria block assimilation of soft wide-angle radiation and provide a halting condition. Mutually hard structures are preserved to the end of clustering, automatically generating a tree of subjet axis candidates. Excellent object identification and kinematic reconstruction for multipronged resonances are realized across more than an order of magnitude in transverse energy. The clustering measure history facilitates high-performance substructure tagging, which we quantify with the aid of supervised machine learning. These properties suggest that SIFT may prove to be a useful tool for the continuing study of jet substructure.

DOI: [10.1103/PhysRevD.108.016005](https://doi.org/10.1103/PhysRevD.108.016005)

I. INTRODUCTION

The collider production of an isolated partonic object bearing uncanceled strong nuclear charge is immediately followed by a frenzied showering of soft and collinear radiation with a complex process of recombination into metastable color-singlet hadronic states. In order to compare theoretical predictions for hard scattering events against experimental observations, it is necessary to systematically reassemble these “jets” of fragmentary debris into a faithful representation of their particle source. The leading clustering algorithm serving this purpose at the Large Hadron Collider (LHC) is anti- k_T [1], which is valued for yielding regular jet shapes that are simple to calibrate.

It can often be the case that clustering is complicated by early decays of a heavy unstable state into multiple hard prongs, e.g., as $(W^+ \rightarrow u\bar{d})$, or $(t \rightarrow W^+ b \rightarrow u\bar{d}b)$. This additional structure can be of great benefit for tagging presence of the heavy initial state. However, identification will be confounded if essential constituents

of distinct prongs either remain uncollected or are merged together into a joint assemblage. Standard practice is to err in the latter direction, via construction of a large-radius jet that is intended to encapsulate all relevant showering products and to subsequently attempt recovery of the lost “substructure” using a separate algorithm such as N subjettiness [2].

The casting of this wide net invariably also sweeps up extraneous low-energy radiation at large angular separations and successful substructure tagging typically hinges on secondarily “grooming” large-radius jets with a technique such as soft drop [3]. Furthermore, the appropriate angular coverage is intrinsically dependent upon the process being investigated, since more boosted parents will tend to produce more collimated beams of children. Accordingly, such methods are conventionally tuned to target specific energy scales for maximal efficacy.

In this paper, we introduce a new jet clustering algorithm called SIFT (scale-invariant filtered tree) that is engineered to avoid losing resolution of structure in the first place. Similar considerations previously motivated development of the exclusive X Cone [4,5] algorithm based on minimization of N jettiness [6]. Most algorithms in common use reference a fixed angular cone size R_0 , inside of which all presented objects will cluster and beyond which all exterior objects will be excluded. We identify this parameter, and the conjugate momentum

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

scale which it imprints, as primary culprits responsible for the ensuing loss of resolution.

In keeping, our proposal rejects any such references to an external scale, while asymptotically recovering successful angular and kinematic behaviors of algorithms in the k_T family. Specifically, the SIFT prioritizes pairing objects that have hierarchically dissimilar transverse momentum scales (i.e., one member is soft *relative* to its partner) and narrow angular separation (i.e., the members are collinear). The form we are lead to by these considerations is quite similar to a prior clustering measure named Geneva [7], although our principal motivation (retention of substructure) differs substantively from those that applied historically.

We additionally resolve a critical fault that otherwise precludes the practical application of radius-free measures such as Geneva, namely the tendency to gather uncorrelated soft radiation at wide angular separations. This is accomplished with a novel filtering and halting prescription that is motivated by soft drop, but defined in terms of the scale-invariant measure itself and applied to each candidate merger during the initial clustering.

The final-state objects retained after algorithm termination are isolated variable-large-radius jets that dynamically bundle decay products of massive resonances in response to the *process* scale. Since mutually hard prongs tend to merge last (this behavior is very different from that of anti- k_T), the end stages of clustering densely encode information indicating the presence of substructure. Accordingly, we introduce the concept of an N -subjett tree, which represents the ensemble of a merged object's projections onto N axes, as directly associated with the clustering progression from N prongs to 1.

We demonstrate that the subjett axes obtained in this manner are effective for applications such as the computation of N subjettiness. Additionally, we confirm that they facilitate accurate and sharply peaked reconstruction of associated mass resonances. Finally, we show that the sequential evolution history of the SIFT measure across mergers is itself an excellent substructure discriminant. As quantified with the aid of a boosted decision tree (BDT), we find that it significantly outperforms the benchmark approach to distinguishing one-, two-, and three-prong events using N -subjettiness ratios, especially in the presence of a large transverse boost.¹

The outline of this work is as follows. Section II presents a master sequential jet clustering algorithm into which SIFT and the k_T prescriptions may be embedded. Sections III–VI define the SIFT algorithm in terms of its scale-invariant measure (with transformation to a coordinate representation), filtering and isolation criteria, and final-state N -subjett tree objects. Section VII visually contrasts clustering

priorities, soft radiation catchment, and halting status against common algorithms. Sections VIII and IX comparatively assesses SIFT's performance in various applications, relating firstly to jet resolution and mass reconstruction, and secondly, to the tagging of structure. Section X addresses computability, infrared, and collinear safety, and $\mathcal{O}(1)$ deviations from recursive safety. Section XI concludes and summarizes. Appendix A provides a pedagogical review of hadron collider coordinates. Appendix B describes available software implementations of the SIFT algorithm along with methods for reproducing current results and plans for integration with the FastJet [10,11] contributions library.

II. THE MASTER ALGORITHM

This section summarizes the global structure of a general sequential jet clustering algorithm, as illustrated in Fig. 1. Examples are provided of how standard algorithms fit into this framework. These examples will be referenced subsequently to motivate the SIFT algorithm and to comparatively assess its performance.

To begin, a pool of N low-level physics objects (e.g., four-vector components of track-assisted calorimeter hits) is populated, typically including reconstructed hadrons as well as photons and light leptons that fail applicable hardness, identification, or isolation criteria.

The main loop then begins by finding the two most-proximal objects (A , B), as defined via minimization of a specified distance measure δ_{AB} over each candidate pairing. For example, the anti- k_T algorithm is a member (with index $n = -1$) of a broader class of algorithms that includes the earlier k_T [12,13] (with $n = +1$) and Cambridge-Aachen [14,15] (with $n = 0$) formulations, corresponding to the measure $\delta_{AB}^{k_T, n}$ defined as

$$\delta_{AB}^{k_T, n} \equiv \min[(p_T^A)^{2n}, (p_T^B)^{2n}] \times \left(\frac{\Delta R_{AB}}{R_0} \right)^2. \quad (1)$$

The quantity $(\Delta R_{AB})^2 \equiv (\Delta\eta)^2 + (\Delta\phi)^2$ expresses “geometric” adjacency as a Cartesian norm square of differences in pseudorapidity η and azimuthal angle ϕ (cf. Appendix A) relative to the maximal cone radius R_0 . The ($n = 0$) scenario prioritizes small values of ΔR_{AB} without reference to the transverse momentum p_T . A positive momentum exponent ($n = +1$) first associates pairs wherein at least one member is very soft, naturally rewinding the chronology of the showering process. A negative momentum exponent ($n = -1$) favors pairs wherein at least one member represents hard radiation directed away from the beamline.

Once an object pair has been selected, there are various ways to handle its members. Broadly, three useful alternatives are available, identified here as clustering, dropping, and isolating. The criteria for distinguishing between these actions are an important part of any algorithm's halting condition. Clustering means that the objects (A , B)

¹For other approaches to scale-invariant resonance tagging, variable-radius jet clustering, and single-pass grooming and subjett finding, see Refs. [8,9].




FIG. 1. Generalized logical flow chart representing a family of sequential jet clustering algorithms defined by a distance measure δ_{AB} , along with a specification for filtering and isolation criteria and/or a target final-state jet count for exclusive clustering.

are merged, usually via a four-vector sum ($p_{AB}^\mu = p_A^\mu + p_B^\mu$), and replaced by this merged object in the pool. Typically, clustering proceeds unless the applicable angular separation is too great and/or the momentum scales of the members are too dissimilar.

For pairings failing this filter, two options remain. Dropping, wherein the softer member is set aside (it may literally be discarded or rather reclassified as a final-state object) while the harder member is returned to the object pool, sensibly applies when momentum scales are hierarchically imbalanced. Conversely, a symmetric

treatment of both members is motivated when momentum scales are similar and isolating involves mutual reclassification as final-state objects. Exclusive algorithms (which guarantee a fixed ending count N_{exc} of jets) typically always cluster, although any construction that reduces the net object count by one unit per iteration is consistent.²

To complete the prior example, conventional implementations of k_T -family clustering do not discard objects or

²A valid example is clustering plus dropping without isolation.

collectively isolate final-state pairs. However, they do singly reclassify all objects with no partners nearer than R_0 in ΔR as final-state jets. This behavior may be conveniently embedded into the master algorithm flow by also allowing each object in the active pool to pair one at a time with the “beam” and associating it with the alternative measure $\delta_{\text{beam}}^{k_T, n} = p_T^{2n}$. The dropping criterion is then adapted to set aside any object for which this beam distance is identified as the global minimum. Note that such objects are indeed guaranteed to have no neighbors inside a centered cone of size R_0 .

The described loop repeats until the number of objects remaining in the active pool reaches a specified threshold. In the context of an exclusive clustering mode, this would correspond to a target count N_{exc} , whereas continuation otherwise simply requires the presence of multiple ($N > 1$) objects. Many clustering algorithms, including SIFT, may be operated in either of these modes. Upon satisfaction of the halting criteria, the list of final-state objects is returned along with a record of the clustering sequence and associated measure history.

The SIFT algorithm will now be systematically developed by providing specific prescriptions for each element of the master algorithm. The scale-invariant measure δ_{AB} is defined in Sec. III and provided with an intuitive geometric form in Sec. IV. The filtering and isolation criteria are itemized in Sec. V. The N -subject tree is introduced in Sec. VI.

III. THE SCALE-INVARIANT MEASURE

This section establishes the SIFT clustering measure δ_{AB} and places it in the context of similar constructs from the literature. Our principal objective is to develop an approach that is intrinsically resilient against loss of substructure in boosted event topologies, i.e., which maintains resolution of collimated radiation associated with distinct partonic precursors in the large- p_T limit.

We identify specification of an angular size parameter R_0 as the primary culprit imposing a conjugate momentum scale dependence on the performance of traditional approaches to jet clustering. In pursuit of a scale-independent algorithm, we require that the clustering measure be free of any such factor. Nevertheless, it is desirable to asymptotically recover angular and kinematic characteristics of existing successful approaches such as anti- k_T [cf. Eq. (1) with $n = -1$], and we thus seek out proxies for its dominant behaviors. Specifically, these include a preference for pairs with close angular proximity and a preference for pairs where one member carries a large transverse boost. For the former purpose, we invoke the mass-square difference, defined as follows:

$$\begin{aligned} \Delta m_{AB}^2 &\equiv (p_A^\mu + p_B^\mu)^2 - m_A^2 - m_B^2 = 2p_A^\mu p_B^\mu \\ &\simeq 2E^A E^B (1 - \cos \Delta\theta_{AB}) \simeq E^A E^B \Delta\theta_{AB}^2. \end{aligned} \quad (2)$$

The property that small mass-square changes correlate with collinearity of decay products was similarly leveraged by the early JADE [16] algorithm. For the latter purpose, we turn to suppression in a denominator by the summed transverse energy-square,

$$\begin{aligned} \sum E_T^2 &\equiv (E_T^A)^2 + (E_T^B)^2 \\ E_T^2 &\equiv p_T^2 + m^2 = E^2 - p_z^2. \end{aligned} \quad (3)$$

The summation plays a role similar to that of the “min” criterion in Eq. (1), in the case that the pair of objects under consideration is very asymmetrically boosted. The choice of E_T^2 over simply p_T^2 prevents a certain type of divergence, as it is possible for the vector quantity $\vec{p}_T$ to cancel during a certain phase of the clustering, but not without generation of mass. All together, the simple prescription for the proposed algorithm is to sequentially cluster indexed objects A and B corresponding to the smallest pairwise value of the following expression:

$$\delta_{AB} \equiv \frac{\Delta m_{AB}^2}{(E_T^A)^2 + (E_T^B)^2}. \quad (4)$$

Since this measure represents our default context, we write it without an explicit superscript label for simplicity. We note that it consists of the dimensionless ratio of a Lorentz invariant and an invariant under longitudinal boosts, and that it manifestly possesses the desired freedom from arbitrary externally-specified scales. For comparison, the JADE clustering measure is

$$\delta_{AB}^{\text{JADE}} \equiv \frac{E_A E_B}{s} \times (1 - \cos \Delta\theta_{AB}). \quad (5)$$

Having been designed for application at a lepton collider, JADE [17] references total energies (E_A, E_B) and spherically symmetric angular separations $\Delta\theta_{AB}$ rather than cylindrical quantities. It also treats all merger candidates as massless; i.e., it uniformly factors ($|\vec{p}| \Rightarrow E$) out of the angular dependence. Although Eq. (5) is dimensionless and avoids referencing a fixed cone size, it is not scale invariant, due to normalization against the global Mandelstam center-of-momentum energy $\sqrt{s}$. This distinction is amplified in a hadron collider context, where the partonic $\sqrt{\hat{s}}$ and laboratory center-of-momentum frames are generally not equivalent.

However, the historical measure most alike to SIFT is Geneva, which was developed as a modification to JADE,

$$\delta_{AB}^{\text{Geneva}} \equiv \frac{8}{9} \frac{E_A E_B}{(E_A + E_B)^2} \times (1 - \cos \Delta\theta_{AB}). \quad (6)$$

Like Eq. (4), Eq. (6) achieves scale invariance by constructing its denominator from dimensionful quantities local to the clustering process. There are also several

differences between the two forms, which are summarized here and explored further in Sec. IV. Variations in the dimensionless normalization are irrelevant.³ The substitution of cylindrical for spherical kinematic quantities is relevant, albeit trivial. Exchanging the squared sum for a sum of squares causes the SIFT measure fall off more sharply than the Geneva measure as the momentum scales of the clustering candidates diverge. But, the most critical distinguishing feature is that SIFT is sensitive to accumulated mass.

Despite core similarities at measure level, the novel filtering and isolation criteria described in Sec. V are responsible for essential divergences in the properties of final-state objects produced by SIFT relative to those from Geneva. In particular, we will show that the N -subjettree introduced in Sec. VI is effective for reconstructing hard objects and tagging the presence of substructure at large transverse boosts without the need for additional de/reclustering or postprocessing.

IV. THE GEOMETRIC FORM

This section presents a transformation of the SIFT clustering measure into the geometric language of coordinate differences. While the underlying physics is invariant under this change of variables, the resulting expression is better suited for developing intuition regarding clustering priorities and how they compare to those of traditional algorithms. It will also be used to motivate and express the filtering and isolation criteria in Sec. V, and to streamline calculations touching on computational safety in Sec. X. Additionally, it is vital for realizing fast numerical implementations of the SIFT algorithm that employ optimized data structures based on geometric coordinate adjacency.

The canonical form of distance measure on a manifold is a sum of bilinear coordinate differentials $dx^i dx^j$ with general coordinate-dependent coefficients g_{ij} . Measures applicable to jet clustering must be integrated, referencing finite coordinate separations Δx^i . The measure expressed in Eq. (4) does not apparently refer to coordinate differences at all, although Eq. (2) provides a hint of how an implicit dependence of this type might manifest. We start from the mass-square difference,

$$\Delta m_{AB}^2 = 2 \times (E^A E^B - p_z^A p_z^B - p_T^A p_T^B \cos \Delta \phi_{AB}). \quad (7)$$

To proceed, recall that Lorentz transformations are generated by hyperbolic “rotation” in rapidity y . In particular, we may boost via matrix multiplication from the transverse frame with $y = 0$, $p_z = 0$, and $E = E_T$ to any longitudinally related frame,

$$\begin{pmatrix} E \\ p_z \end{pmatrix} = \begin{pmatrix} \cosh y & \sinh y \\ \sinh y & \cosh y \end{pmatrix} \begin{pmatrix} E_T \\ 0 \end{pmatrix} = \begin{pmatrix} E_T \cosh y \\ E_T \sinh y \end{pmatrix}. \quad (8)$$

This may be used to reduce Eq. (7), using the standard hyperbolic difference identity,

$$\begin{aligned} E^A E^B - p_z^A p_z^B &= E_T^A E_T^B \times (\cosh y^A \cosh y^B - \sinh y^A \sinh y^B) \\ &= E_T^A E_T^B \times \cosh \Delta y_{AB}. \end{aligned} \quad (9)$$

The transverse energy will factor perfectly out of the mass-square difference if the constituent four-vectors are individually massless, i.e., if $(m_A = m_B = 0)$. Otherwise, there are residual coefficients (ξ_A, ξ_B) defined as follows:

$$\begin{aligned} \Delta m_{AB}^2 &= 2 E_T^A E_T^B \times (\cosh \Delta y_{AB} - \xi^A \xi^B \cos \Delta \phi_{AB}) \\ \xi &\equiv \frac{p_T}{E_T} = \left(1 - \frac{m^2}{E_T^2}\right)^{+1/2} = \left(1 + \frac{m^2}{p_T^2}\right)^{-1/2}. \end{aligned} \quad (10)$$

The role of the ξ in Eq. (10) is to function as a “lever arm” deemphasizing azimuthal differences in the non-relativistic limit and at low p_T . The precise relationship between Δm_{AB}^2 and the angular separation ΔR_{AB}^2 can also now be readily established, referencing Taylor expansions for cosine and hyperbolic cosine in the limit of small angular separations with $(p_T \ll m)$,

$$\begin{aligned} \Delta \tilde{R}_{AB}^2 &\equiv \frac{\Delta m_{AB}^2}{E_T^A E_T^B} \\ &= 2 \times (\cosh \Delta y_{AB} - \xi^A \xi^B \cos \Delta \phi_{AB}) \\ &\simeq \Delta \eta_{AB}^2 + \Delta \phi_{AB}^2 \equiv \Delta R_{AB}^2. \end{aligned} \quad (11)$$

The indicated correspondence becomes increasingly exact as one approaches the collinear ($\Delta R^2 \ll 1$) and massless ($\Delta y \Rightarrow \Delta \eta$, $\xi \Rightarrow 1$) limits. The fact that $\cosh \Delta y$ is unbounded and has purely positive series coefficients, whereas $\cosh \Delta \phi$ is bounded and its terms are of alternating sign, implies that $\Delta \tilde{R}^2$ is more sensitive to separations in rapidity than separations in azimuth.

The ratio $(\Delta \tilde{R}^2 \div \Delta R^2)$ is explored graphically in Fig. 2 as a function of $\Delta \eta$ and $\Delta \phi$, for various values of (m/p_T) . For purposes of illustration, deviations from $(\eta = 0)$ are applied symmetrically to the candidate object pair, and the quoted mass ratio applies equivalently to both objects. Black contours indicate unity, bisecting regions of neutral bias, which are colored in blue. Gray contours are spaced at increments of 0.1 in the base-10 logarithm, and regions where SIFT exhibits enhancement (suppression) of clustering are colored green (red).

In the massless limit (left-hand panel) a near symmetry persists under $(\Delta \eta \Leftrightarrow \Delta \phi)$, with corrections from subleading terms as described previously. When m approaches p_T

³Geneva was introduced with a coefficient of 8/9 to match the normalization of JADE in the limit of three hard prongs.




FIG. 2. The factor $\Delta\tilde{R}_{AB}^2$ is compared to the traditional angular measure ΔR_{AB}^2 , as a function of $\Delta\eta$ and $\Delta\phi$, and for various values of (m/p_T) . Regions where the SIFT measure enhances (suppresses) clustering are shown in green (red). Blue indicates similar behavior, with black contour placed at unity. Gray contours mark increments of 0.1 in $\log_{10}(\Delta\tilde{R}_{AB}^2 / \Delta R_{AB}^2)$.

(center panel), the pseudorapidity versus azimuth symmetry is meaningfully broken at large ΔR , and a strong aversion emerges to the clustering of massive states at small ΔR . The latter effect dominates for highly nonrelativistic objects (right-hand panel), and extends to larger separations, such that the distinction between $\Delta\eta$ and $\Delta\phi$ is again washed out. SIFT binds objects approaching ($\Delta R \sim \pi$) significantly more tightly than the k_T algorithms in this limit. Intuition for that crossover in sign can be garnered from the leading terms in the multivariate expansion shown in the following, as developed from Eqs. (10), (11), and (A5):

$$\Delta\tilde{R}_{AB}^2 \Rightarrow \Delta R_{AB}^2 + \left\{ 1 - \frac{\Delta R_{AB}^2}{2} \right\} \times \left\{ \left(\frac{m_A}{p_T^A} \right)^2 + \left(\frac{m_B}{p_T^B} \right)^2 \right\} + \dots \quad (12)$$

We turn attention next to the denominator from Eq. (4), defining a new quantity ϵ^{AB} in conjunction with the transverse energy product factored out of Δm_{AB}^2 ,

$$\epsilon^{AB} \equiv \frac{E_T^A E_T^B}{(E_T^A)^2 + (E_T^B)^2} = \left\{ \left(\frac{E_T^A}{E_T^B} \right) + \left(\frac{E_T^B}{E_T^A} \right) \right\}^{-1}. \quad (13)$$

This expression has a symmetry under the transformation ($\alpha \equiv E_T^A/E_T^B \Rightarrow \alpha^{-1}$). It is maximized at $\alpha = 1$, where ($\epsilon^{AB} \Rightarrow 1/2$), and minimized at $\alpha = (0, +\infty)$, where ($\epsilon^{AB} \Rightarrow 0$). The response is balanced by a change of variables $\Delta u = \ln \alpha$, i.e., $\alpha = e^{\Delta u}$, where the logarithm converts ratios into differences,

$$u \equiv \ln(E_T/[\text{GeV}]) \\ \epsilon^{AB} = (e^{+\Delta u_{AB}} + e^{-\Delta u_{AB}})^{-1} = (2 \cosh \Delta u_{AB})^{-1}. \quad (14)$$

Putting everything together, we arrive at a formulation of the measure from Eq. (4) that is expressed almost entirely in

terms of coordinate differences of the rapidities, azimuths, and log-transverse energies, excepting the coefficients ξ from Eq. (10), which depend on the ratios m/p_T . Of course, it is possible to adopt a reduction of the measure where ($\xi \Rightarrow 1$) by fiat, which is equivalent to taking a massless limit in the fashion of JADE and Geneva, as described in Sec. III,

$$\delta_{AB} = \epsilon^{AB} \times \Delta\tilde{R}_{AB}^2 \\ = \frac{\cosh \Delta y_{AB} - \xi^A \xi^B \cos \Delta\phi_{AB}}{\cosh \Delta u_{AB}}. \quad (15)$$

The scale-invariance of Eq. (15) is explicit, in two regards. By construction, there is no reference to an external angular cutoff R_0 . In addition, the fact that transverse energies are referenced only via ratios, and never in absolute terms, is emergent. The measure is additionally observed to smoothly blend attributes of k_T and anti- k_T jet finding, insomuch as the former prioritizes clustering when one member of a pair is soft, the latter when one member of a pair is hard, and SIFT when the transverse energies are logarithmically disparate.

This behavior is illustrated in Fig. 3, with $2 \times \epsilon^{AB}$ plotted in blue as a function of ($\alpha \equiv E_T^A/E_T^B$). For comparison, the analogous momentum-dependent factor for Geneva from Eq. (6) is shown in green on the same axes, taking ($E \Rightarrow E_T$) and normalizing to unity at ($\alpha = 1$). Like SIFT, Geneva is insensitive to variation of the absolute event scale ($\beta \equiv E_T^A E_T^B / E_0^2$). However, the extra cross term appearing in its denominator produces heavier tails when the scale ratio α is unbalanced. The cusped red region similarly represents the k_T and anti- k_T measures. Specifically, the following expression is proportional to Eq. (1) for ($n = \pm 1$) in the massless limit:

$$\delta_{AB}^{k_T, n} \propto \left(\frac{E_T^A E_T^B}{E_0^2} \right)^n \times \min \left[\frac{E_T^A}{E_T^B}, \frac{E_T^B}{E_T^A} \right]. \quad (16)$$




FIG. 3. The factor $2 \times \epsilon^{AB}$ is log symmetric in the ratio E_T^A/E_T^B of transverse energies, becoming small whenever candidate scales are hierarchically dissimilar. SIFT (blue) is compared against analogous behaviors for the Geneva (green) and k_T -family (red) algorithms, the latter at $(\beta \equiv E_T^A E_T^B / E_0^2 \Rightarrow 1)$. Gray contours illustrate scaling of the k_T measures with a power $n = \pm 1$ of $(\beta \Rightarrow 1/10, 1/5, 1/2, 2, 5, 10)$.

The selected normalization agrees with $2 \times \epsilon^{AB}$ in the further limits $(\beta \Rightarrow 1)$ and $(\alpha \Rightarrow 1)$, where E_0 is an arbitrary constant reference energy. The distinction between k_T and anti- k_T clustering amounts to an enhancement versus suppression by the product (squared geometric mean) of transverse momenta. This is illustrated with the gray contours in Fig. 3, which rescale by $\beta = (1/10, 1/5, 1/2, 2, 5, 10)$ from inner lower to outer upper for $(n = +1)$, or in the reverse order for $(n = -1)$. The β -invariant case $(n = 0)$ is also potentially of interest, but it is a new construction that is not to be confused with the Cambridge-Aachen algorithm, which has no energy dependence at all.

We conclude this section with a transformation that identifies the coordinate u introduced in Eq. (14) as a sort of “dual” to the rapidity y from Eq. (A5). The log-transverse momentum $\ln(p_T/[\text{GeV}])$ is similarly linked to pseudorapidity η in the massless limit,

$$\begin{aligned}\Omega_{\pm} &\equiv \ln\left(\frac{E \pm p_z}{[\text{GeV}]}\right) \\ \frac{\Omega_+ + \Omega_-}{2} &= \frac{1}{2} \ln\left(\frac{E^2 - p_z^2}{[\text{GeV}]^2}\right) = u \\ \frac{\Omega_+ - \Omega_-}{2} &= \frac{1}{2} \ln\left(\frac{E + p_z}{E - p_z}\right) = y.\end{aligned}\quad (17)$$

V. FILTERING, ISOLATION, AND HALTING

This section establishes the SIFT filtering and isolation criteria, which are used in conjunction to formulate a suitable halting condition for the nonexclusive clustering mode. Direct integration of a grooming stage effectively rejects stray radiation and pileup. SIFT’s behavior will be

visualized with and without filtering in Sec. VII, and compared against each k_T -family algorithm in the presence of a soft “ghost” radiation background.⁴

In conjunction, the two factorized terms in Eq. (15) ensure that clustering prioritizes the merger of object pairs that have a hierarchically soft member and/or that are geometrically collinear, mimicking fundamental poles in the matrix element for QCD showering. This behavior is illustrated by the “phase diagram” in Fig. 4, where the product of horizontal “ x ” and vertical “ $y(x)$ ” coordinates on that plane is equal to δ_{AB} . Gray “ $y(x) = 1/x$ ” contours trace constant values of the measure, equal to (.002, .005, .01, .02, .05, .1, .2, .5), with minimal values gathered toward the lower left. As a consequence, SIFT successfully preserves mutually hard structures with tight angular adjacency, maintaining their resolution as distinct objects up to the final stages of clustering.⁵

However, iterative application of the SIFT (or Geneva) measure does not offer an immediately apparent halting mechanism, and it will ultimately consume any presented objects into a single all-encompassing jet if left to run. These measures are additionally prone to sweeping up uncorrelated soft radiation onto highly boosted partners at wide angular separation. The solution to both problems turns out to be related. For inspiration, we turn to the soft drop procedure, wherein a candidate jet is recursively declustered and the softer of separated constituents is discarded until the following criterion is satisfied:

⁴The ability to effectively filter soft radiation is suggestive of resiliency to pileup. Although beyond the scope of this work, it will be interesting to test how SIFT’s performance scales as additional minimum-bias collisions are overlaid.

⁵Mutually soft pairings tend not to occur, since such objects are typically gathered up by a harder partner at an early stage.




FIG. 4. Phase diagram for the separation of object merging, filtering, and isolation responses.

$$\frac{\min(p_T^A, p_T^B)}{p_T^A + p_T^B} > z_{\text{cut}} \left(\frac{\Delta R_{AB}}{R_0} \right)^\beta. \quad (18)$$

The dimensionless z_{cut} coefficient is typically $\mathcal{O}(0.1)$. The exponent β can vary for different applications, although we focus here on $\beta = 2$. We first attempt to recast the elements of Eq. (18) into expressions with asymptotically similar behavior that adopt the vocabulary of Eq. (4). The factor $\Delta \tilde{R}_{AB}^2$ can be carried over directly. Likewise, ϵ^{AB} behaves similarly to a minimized ratio of transverse energies, with the advantage of analyticity,

$$\begin{aligned} \epsilon^{AB} &= \left\{ \min \left(\frac{E_T^A}{E_T^B} \right) + \max \left(\frac{E_T^A}{E_T^B} \right) \right\}^{-1} \\ &\simeq \left\{ \max \left(\frac{E_T^A}{E_T^B} \right) \right\}^{-1} \\ &= \min \left(\frac{E_T^A}{E_T^B} \right) \simeq \frac{\min(E_T^A, E_T^B)}{E_T^A + E_T^B}. \end{aligned} \quad (19)$$

Paired factors of 2 previously emerged “naturally” in Eqs. (11) and (14), and we group them here in a way that could be interpreted as setting “ $z_{\text{cut}} = 1/4$ ” in the context of a large-radius ($R_0 = 1$) jet. Putting all of this together, the suggested analog to Eq. (18) is shown in the following:

$$\text{Cluster: } \frac{\Delta \tilde{R}_{AB}^2}{2} < \{(2\epsilon^{AB}) \leq 1\}. \quad (20)$$

An important distinction from conventional usage is that this protocol is to be applied during the initial clustering cycle itself, at the point of each candidate merger. This is similar in spirit to recursive soft drop [18].

Intuition may be garnered by turning again to Fig. 4, where clustering consistent with the Eq. (20) prescription is observed to occur only in the upper-left (green) region of

the phase diagram, above the “ $y(x) = x$ ” diagonal cross-cutting contours of the measure. The upper bound on $(2\epsilon^{AB})$ precludes clustering if $(\Delta \tilde{R}_{AB} \geq \sqrt{2})$. However, this angular threshold is *dynamic*, and the capturable cone diminishes in area with increasing imbalance of the transverse scales. This condition may be recast into a limit on the clustering measure from Eq. (15),

$$\text{Cluster: } \delta_{AB} < \{(2\epsilon^{AB})^2 \leq 1\}. \quad (21)$$

The question now is what becomes of those objects rejected by the Eq. (21) filter. Specifically, are they discarded or are they classified for retention as isolated final-state jets? One possible solution is to simply determine this based on the magnitude of the transverse energy but that runs somewhat counter to current objectives. Seeking a simple scale-invariant criterion for selecting between isolation and rejection, we observe that there are distinct two ways in which Eq. (20) may fail. Namely, the angular opening may be too wide, and/or the transverse scales may be too hierarchically separated. This is illustrated again by Fig. 4, wherein one may exit the clustering region (green) by moving rightward (wider angular separation) or downward (more scale disparity).

If the former cause is dominant, e.g., if $(\Delta \tilde{R}_{AB}^2 \gg 1)$ and $(\epsilon^{AB} \sim 1)$ such that both members are on equal footing and $(\delta_{AB} \gg 1)$, then collective isolation as a pair of distinct final states is appropriate. Conversely, if the latter cause primarily applies, e.g., if $(\epsilon^{AB} \ll 1)$ and $(\Delta \tilde{R}_{AB}^2 \sim 1)$ such that $(\delta_{AB} \ll 1)$, then the pertinent action is to asymmetrically set aside just the softer candidate.⁶ These scenarios may be quantitatively distinguished in a manner that

⁶Although we treat “drop” here as a synonym for “discard,” a useful alternative (cf. Sec. II) is single reclassification as a final state, since softness is not guaranteed in an absolute sense.

generates a balance with and continuation of Eq. (20), as follows:

$$\begin{aligned} \text{Drop: } \{(2e^{AB}) \leq 1\} &\leq \frac{\Delta\tilde{R}_{AB}^2}{2} < \{1 \leq (2e^{AB})^{-1}\} \\ \text{Isolate: } \{1 \leq (2e^{AB})^{-1}\} &\leq \frac{\Delta\tilde{R}_{AB}^2}{2}. \end{aligned} \quad (22)$$

Each criterion may also be recast in terms of the clustering measure, with isolation always and only indicated against an absolute reference value (1) of the SIFT measure δ_{AB} that heralds a substantial mass gap. Observe that the onset of object isolation necessarily culminates in global algorithmic halting since all residual pairings must correspond to larger values of δ_{AB} ,

$$\begin{aligned} \text{Drop: } \{(2e^{AB})^2 \leq 1\} &\leq \delta_{AB} < \{1\} \\ \text{Isolate: } \{1\} &\leq \delta_{AB}. \end{aligned} \quad (23)$$

Figure 4 again provides visual intuition, where isolation occurs in the upper-right (blue) regions for values of the measure above unity, and soft wide radiation is dropped in the lower-central (red) regions. The clustering and isolation regions are fully separated from each other, making contact only at the zero-area “triple point” with $(2e^{AB} = 1, \Delta\tilde{R}_{AB}^2/2 = 1)$. We will provide additional support from simulation for triggering isolation at a fixed $\mathcal{O}(1)$ value of the measure in Sec. IX.

Note that nonrelativistic or beamlike objects at very large rapidity that approach the $(p_T \Rightarrow 0, m \neq 0)$ limit will never cluster, since vanishing of the lever arm ($\xi = 0$) in Eq. (10) implies via Eq. (11) that $(\Delta\tilde{R}_{AB}^2/2 \geq 1)$. Also, objects of equivalent transverse energy with angular separation $(\Delta\tilde{R}_{AB} \geq \sqrt{2})$ will always be marked for isolation, although this radius is again *dynamic*. The phase gap between clustering and isolation opens up further when scales are mismatched, and the effective angular square separation exceeds its traditional counterpart as mass is accumulated during clustering [cf. Eq. (12)].

VI. THE N -SUBJET TREE

This section concludes development of the SIFT algorithm by formalizing the concept of an N -subjettree. This data structure records the clustering history of each isolated final-state object from N prongs down to 1, along with the associated value of the measure at each merger.

Final-state jets defined according to the global halting condition outlined in Sec. V may still bundle multiple hard structured prongs, since isolation requires a minimal separation of $(\Delta\tilde{R}_{AB} \geq \sqrt{2})$. Within each quarantined partition, SIFT reverts to its natural form, as an exclusive clustering algorithm. So, for example, it might be that reconstruction of a doubly hadronic $t\bar{t}$ event would isolate the pair of top-quark remnants but merge each bottom with

products from the associated W . This is a favorable outcome, amounting to the identification of variable large-radius jets. But, the question of how to identify and recover the optimal partition of each such object into N subjets remains.

One could consider formulating a local halting condition that would block the further assimilation of hard prongs within each large-radius jet once some threshold were met. However, such objects are only defined in our prescription *after* having merged to exhaustion. More precisely, a number of candidate large-radius jets may accumulate objects in parallel as clustering progresses, and each will have secured a unique “ $N = 1$ ” configuration at the moment of its isolation.

As such, the only available course of action appears to be proceeding with these mergers, even potentially in the presence of substructure. However, the fact that the SIFT measure tends to preserve mutually hard features until the final stages of clustering suggests that suitable proxies for the partonic event axes may be generated automatically as a product of this sequential transition through all possible subjet counts, especially during the last few mergers. We refer to the superposed *ensemble* of projections onto $(N = \dots, 3, 2, 1)$ prongs, i.e., the history of residually distinct four-vectors at each level of the clustering flow, as an N -subjettree.

In fact, it seems that interrupting the final stages of clustering would amount to a substantial information forfeiture. Specifically, the merger of axis candidates that are not collinear or relatively soft imprints a sharp discontinuity on the measure, which operates in this context like a mass-drop tagger [19] to flag the presence of substructure. Additionally, the described procedure generates a basis of groomed axes that are directly suitable for the computation of observables such as N subjettness. In this sense, the best way to establish that a pair of constituents within a large-radius jet should be kept apart may be to go ahead and join them, yet to remember what has been joined and at which value of δ_{AB} .

In contrast to conventional methods for substructure recovery that involve de and reclustering according to a variety of disjoint prescriptions, the finding of N -subjettrees representing a compound scattering event occurs in conjunction with the filtering of stray radiation and generation of substructure observables during a single unified operational phase. The performance of this approach for kinematic reconstruction and tagging hard event prongs will be comparatively assessed in Secs. VIII and IX.

VII. COMPARISON OF ALGORITHMS

This section provides a visual comparison of merging priorities and final states for the SIFT and k_T -family algorithms. The images presented here are still frames extracted from full-motion video simulations of each clustering sequence. These films are provided as ancillary




FIG. 5. Upper: A simulated LHC scattering event with $\sqrt{s} = 14$ TeV is visualized at the partonic level. Top quark pair production $pp \rightarrow t\bar{t}$ (red, red) is followed by the decays $t \rightarrow W^+ b$ (green, black) and $W^+ \rightarrow u\bar{d}$ (black, black), plus conjugates. A transverse boost of $p_T \simeq 800$ GeV for each top quark produces narrow collimation of decay products. Lower: Generator-level radiation deposits resulting from showering, hadronization, and decay of the partonic event appearing in the upper frame.

files with the source package for this paper on the arXiv and may also be viewed on YouTube [20]. Frames are generated for every 25th clustering action, as well as each of the initial and final 25 actions. The *Mathematica* notebook used to generate these films is described in Appendix B and maintained with the AEACuS [21] package on GitHub.

The visualized event⁷ comes from a simulation in MadGraph/MadEvent [22] of top quark pair production at the

⁷This event is expected to be reasonably representative, being the first member of its Monte Carlo sample.

14 TeV LHC with fully hadronic decays. The scalar sum H_T of transverse momentum is approximately 1.6 TeV at the partonic level. This large boost results in narrow collimation of the three hard prongs (quarks) on either side of the event, as depicted in the upper frame of Fig. 5. In-plane axes represent the pseudorapidity η and azimuth ϕ , with a cell width ($\Delta R \simeq 0.1$) approximating the resolution of a modern hadronic calorimeter. The height of each block is proportional to the log-transverse momentum $\log_{10}(p_T/[\text{GeV}])$ it contains.

The event is passed through through PYTHIA8 [23] for showering and hadronization, and through DELPHES [24]




FIG. 6. Frames representing sequential clustering of the Fig. 5 event using the exclusive $N_{\text{exc}} = 1$ SIFT algorithm without the associated filtering or isolation criteria (cf. Fig. 10 for nonexclusive clustering with both criteria enabled). Upper: Mutually hard prongs with narrow angular separation remain unmerged up to the final stages of clustering. However, hard objects are likely to sweep up soft radiation at wide angles. Lower: An image of the initial pair production is reconstructed just prior to termination. In the absence of a supplementary halting criterion, these structures will subsequently merge to completion, accompanied by a large discontinuity in the measure δ_{AB} .

for fast detector simulation. Detector effects are bypassed in the current context, which starts with unclustered generator-level (PYTHIA8) jets and nonisolated photons/leptons extracted from the DELPHES event record by AEACuS, but they will be included for most of the analysis in Secs. VIII and IX. This initial state is depicted in the lower frame of Fig. 5, which exhibits two dense clusters of radiation that are clearly associated with the partonic event, as well as several offset deposits having less immediate origins in the underlying event or initial state. Prior to clustering, a sheet

of ultrasoft ghost radiation is distributed across the angular field in order to highlight differences between the catchment area [25] of each algorithm.

We begin with an example of exclusive ($N_{\text{exc}} = 1$) clustering ordered by the SIFT measure from Eq. (4), but without application of the filtering and isolation criteria described in Sec. V. Film A clearly exhibits both of the previously identified pathologies, opening with a sweep of soft-wide radiation by harder partners (as visualized with regions of matching coloration) and closing with the contraction of




FIG. 7. Frames representing sequential clustering of the Fig. 5 event using the anti- k_T algorithm. Upper: Priority is given to the hardest radiation, which immediately captures surrounding territory. Hard substructure at angular scales smaller than the clustering radius will be washed out rapidly. Lower: The final state is characterized by regular jet shapes with uniform expected areas.

hard-wide structures into a single surviving object. However, the described success is manifest in between, *vis-à-vis* mutual preservation of narrowly bundled hard prongs until the end stages of clustering. In particular, the upper frame of Fig. 6 features a pair of triplets ($\delta_{AB} \simeq 0.04$) that fairly approximate their collinear antecedents despite bearing wide catchments. Subsequently, this substructure collapses into an image of the original pair production, as depicted at ($\delta_{AB} \simeq 0.3$) in the lower frame of Fig. 6. Nothing further occurs until ($\delta_{AB} \gtrsim 1.0$), beyond which residual structures begin to merge and migrate in unphysical ways.

For comparison, we process the same event using the anti- k_T algorithm at ($R_0 = 0.5$). Film B shows how early

activity is dominated by the hardest radiative seeds, which promptly capture all available territory up to the stipulated radial boundary. In particular, any substructure that is narrower than R_0 will be rapidly erased, as illustrated in the upper frame of Fig. 7. The subsequent stages of clustering are of lesser interest, being progressively occupied with softer seeds gathering up yet softer unclaimed scraps. At termination, the lower frame of Fig. 7 exhibits the regular cone shapes with uniform catchment areas that are a hallmark of anti- k_T . This property is linked to the anchoring of new cones on hard prongs that are less vulnerable to angular drift. It is favored by experimentalists for facilitating




FIG. 8. Frames representing sequential clustering of the Fig. 5 event using the k_T algorithm. Upper: Priority is given to the softest radiation, resulting in the growth of dispersed catchments having a correlation length that increases in time. Mutually hard substructures are preserved until the last stages of clustering. Lower: The final state is characterized by irregular jet shapes with unpredictable areas.

calibration of jet energy scales and subtraction of soft pileup radiation.

Proceeding, we repeat the prior exercise using the k_T algorithm at ($R_0 = 0.5$). In contrast to anti- k_T , clustering is driven here by the softest seeds. Film C demonstrates the emergence of a fine grain structure in the association pattern of objects from adjacent regions that grows in size as the algorithm progresses. Unlike SIFT, which preferentially binds soft radiation to a hard partner, mutually soft objects without a strong physical correlation are likely to pair in this case. Since summing geometrically adjacent partners tends increase p_T , merged objects become less immediately attractive to the measure. As a result, activity

is dispersed widely across the plane, and attention jumps rapidly from one location to the next. However, the combination of mutually hard prongs is actively deferred, which causes collimated substructures to be preserved, as shown in the upper frame of Fig. 8. In contrast to SIFT, this hardness criterion is absolute, rather than relative. Ultimately, structures more adjacent than the fixed angular cutoff R_0 will still be absorbed. Jet centers are likely to drift substantially, and associated catchment shapes are thus highly irregular, as shown in the lower frame of Fig. 8.

Similarly, we also cluster using the Cambridge Aachen algorithm at ($R_0 = 0.5$). Pairings are driven here solely by angular proximity, and Film D shows an associated growth




FIG. 9. Frames representing sequential clustering of the Fig. 5 event using the Cambridge-Aachen algorithm. Upper: Priority is given only the angular proximity, without reference to the momentum scale. Substructure is preserved only until the grain size eclipses its angular scale. Lower: The final state is characterized by irregular jet shapes with unpredictable areas.

of grain size that is like that of the k_T algorithm. The banded sequencing is simply an artifact of the way we disperse ghost jets, randomizing p_T but regularizing placement on the grid. In contrast, mutually hard substructures are not specifically protected and will last only until the correlation length catches up to their separation, as shown in the upper frame of Fig. 9. As before, the angular cutoff R_0 limits resolution of structure. Likewise, jet drift leads to irregular catchment shapes, as shown in the lower frame of Fig. 9.

We conclude this section with a reapplication of the SIFT algorithm, enabling the filtering and isolation criteria from Sec. V. Film E demonstrates that the soft ghost radiation is

still targeted first, but it is now efficiently discarded (as visualized with dark gray) rather than clustered, suggesting resiliency to soft pileup. Hard substructures are resolved without accumulating stray radiation, as shown in the upper frame of Fig. 10. This helps to stabilize reconstructed jet kinematics relative to the parton-level event. Objects decayed and showered from the pair of opposite-hemisphere top quarks are fully isolated from each other in the final state, as shown in the lower frame of Fig. 10. Each such object selectively associates constituents within a “fuzzy” scale-dependent catchment boundary. Nevertheless, it may still be possible to establish an “effective” jet radius by integrating the pileup distribution function up to the maximal radius




FIG. 10. Frames representing sequential clustering of the Fig. 5 event using the SIFT algorithm with the application of filtering and isolation criteria. Upper: Initial activity is dominated by the rejection of soft-wide radiation that is paired with a hard prong by the measure but fails the filtering criterion. Mutually hard substructures are resolved without contamination from stray radiation. Lower: The isolation criterion triggers halting before distinct objects associated with the initial pair production would be merged.

($\Delta\tilde{R}_{AB} < \sqrt{2}$). In any case, the traditional approach to pileup subtraction has been somewhat superseded by the emergence of techniques for event-by-event, and per-particle pileup estimation like PUPPI [26].

VIII. RESOLUTION AND RECONSTRUCTION

This section characterizes SIFT's angular and energetic response functions for the resolution of hard monojets and tests the reconstruction of collimated di- and trijet systems associated with a massive resonance. The best performance is achieved for large transverse boosts.

We generate Monte Carlo collider data modeling the $\sqrt{s} = 14$ TeV LHC using MadGraph/MadEvent, PYTHIA8, and DELPHES as before. Clean ($N = 1, 2, 3$) prong samples are obtained by simulating the processes ($pp \Rightarrow jZ \Rightarrow j + \nu\bar{\nu}$), ($pp \Rightarrow W^\pm Z \Rightarrow jj + \nu\bar{\nu}$), and ($pp \Rightarrow tW^- \Rightarrow jjj + \bar{\nu}\ell^-$) plus conjugate, respectively. In the latter case, an angular isolation cone with ($\Delta R = 0.5$) is placed around the visible lepton. Hard partonic objects are required to carry a minimal transverse momentum ($p_T \geq 25$ GeV) and be inside ($|\eta| \leq 3.0$). No restrictions are placed on the angular separation of decay products. Jets consist of gluons and/or light first-generation quarks (u, d), as well as b




FIG. 11. Top: Distribution $\mathcal{R}_B^A$ of reconstructed jet energy responses with detector effects relative to the partonic truth level at various transverse boosts. Bottom: Distribution ΔR_B^A of reconstructed jet angular responses with detector effects relative to the partonic truth level at various transverse boosts. Left-hand panels represent the leading filtered and isolated SIFT jet, while right-hand panels represent the leading ($R_0 = 1$) large-radius soft drop jet reported by DELPHES. No calibration of jet energy scales is attempted for either category.

quarks where required by a third-generation process. In order to represent a wide range of event scales, we tranche in the transverse momentum (vector sum magnitude) of the hadronic system, considering six log-spaced intervals $p_T = (100, 200, 400, 800, 1600, 3200)$ GeV $\pm 5\%$ and giving attention primarily to the inner four.

Clustering is disabled at the detector simulation level by setting the jet radius R_0 and aggregate p_T threshold to very small values. We retain the default DELPHES efficiencies for tracks and calorimeter deposits (including p_T thresholds on low-level detector objects), along with cell specifications and smearing (resolution) effects in the latter case. Jet energy scale corrections are turned off (set to 1.0) since these are calibrated strictly for application to fully reconstructed (clustered) objects. For purposes of comparison and validation, we also extract information from DELPHES regarding the leading large-radius jet ($R_0 = 1$), which is processed by trimming [27], pruning [28], and applying soft drop.

Event analysis (including clustering) and computation of observables are implemented with AEACuS (cf. Appendix B). We begin by preclustering detector-

level objects with anti- k_T at ($R_0 = 0.01$) to roughly mimic a characteristic track-assisted calorimeter resolution at the LHC. The isolation and filtering criteria described in Sec. V are then used in conjunction to select the subset of detector-level object candidates retained for analysis. Specifically, our procedure is equivalent to keeping members gathered by the hardest isolated N -subjettree that survive filtering all the way down to the final merger. All histograms are generated with RHADAManTHUS [21], using Matplotlib [29] on the back end.

We begin by evaluating the fidelity of the variable large-radius SIFT jet's directional and scale reconstruction in the context of the monojet sample. Respectively, the upper-left and lower-left panels of Fig. 11 show distributions of the energy response $\mathcal{R}_B^A \equiv (p_T^A/p_T^B - 1)$ and angular response ΔR_B^A relative to the original truth-level (MadGraph) partonic jet at various transverse boosts. The corresponding right-hand panels feature the same two distributions for the leading large-radius jet identified by DELPHES. The vanishing tail of events for which the SIFT or DELPHES jet fails ($\Delta R \leq 0.5$) relative to the partonic sum are vetoed here and throughout.

TABLE I. Detector-level jet energy responses $\mathcal{R}_B^A$ and angular responses ΔR_B^A with associated resolutions $\sigma_{\mathcal{R}}$ and $\sigma_{\Delta R}$.

$p_T^{\text{GeV} \pm 5\%}$	$\langle \mathcal{R}_{\text{Truth}}^{\text{SIFT}} \rangle$	$\sigma_{\mathcal{R}}^{\text{SIFT}}$	$\langle \mathcal{R}_{\text{Truth}}^{\text{DELPHES}} \rangle$	$\sigma_{\mathcal{R}}^{\text{DELPHES}}$	$\langle \Delta \mathcal{R}_{\text{Truth}}^{\text{SIFT}} \rangle$	$\sigma_{\Delta R}^{\text{SIFT}}$	$\langle \Delta \mathcal{R}_{\text{Truth}}^{\text{DELPHES}} \rangle$	$\sigma_{\Delta R}^{\text{DELPHES}}$
100	-0.009	0.17	+0.087	0.17	0.17	0.12	0.17	0.12
200	-0.046	0.16	+0.026	0.16	0.15	0.12	0.15	0.12
400	-0.059	0.15	-0.002	0.15	0.13	0.11	0.13	0.11
800	-0.071	0.14	-0.024	0.14	0.10	0.10	0.10	0.10
1600	-0.081	0.13	-0.042	0.13	0.08	0.09	0.08	0.09
3200	-0.089	0.12	-0.058	0.12	0.05	0.06	0.05	0.06

Central values and associated widths (standard deviations) are provided for the $\mathcal{R}_B^A$ and ΔR_B^A in Table I. The SIFT variable-radius jet energy response is very regular, systematically under-estimating the momentum of hard objects by about 6%. The DELPHES jet energy response shows more drift, transitioning from positive values for soft objects to negative values for hard objects. Widths of the two distributions are indistinguishable, with fluctuations amounting to about 15% in both cases, narrowing slightly at larger boosts. It is anticipated from these observations that an energy calibration of SIFT jets would be relatively

straightforward, using standard techniques. Angular performance of the two methodologies is identical, with typical offsets and deviations both near one-tenth of a radian (but less for hard objects and more for soft objects).

For comparison, we repeat this analysis in Fig. 12 and Table II, using generator-level (PYTHIA8) objects without detector effects and suppressing the emission of initial-state radiation. The distributions are substantially narrower in all cases, with widths around a half or a third of the prior reference values. The energy response is affected by both idealizations, but more so by the elimination of detector




FIG. 12. Top: Distribution $\mathcal{R}_B^A$ of reconstructed jet energy responses, at generator level and without initial-state radiation, relative to the partonic truth level at various transverse boosts. Bottom: Distribution ΔR_B^A of reconstructed jet angular responses, at generator level and without initial-state radiation, relative to the partonic truth level at various transverse boosts. Left-hand panels represent the leading filtered and isolated SIFT jet, while right-hand panels represent the leading ($R_0 = 1$) large-radius soft drop jet reported by DELPHES. No calibration of jet energy scales is attempted for either category.


TABLE II. Zero-ISR generator-level jet energy responses $\mathcal{R}_B^A$, and angular responses ΔR_B^A with associated resolutions $\sigma_{\mathcal{R}}$ and $\sigma_{\Delta R}$.

$p_T^{\text{GeV} \pm 5\%}$	$\langle \mathcal{R}_{\text{Truth}}^{\text{SIFT}} \rangle$	$\sigma_{\mathcal{R}}^{\text{SIFT}}$	$\langle \mathcal{R}_{\text{Truth}}^{\text{DELPHES}} \rangle$	$\sigma_{\mathcal{R}}^{\text{DELPHES}}$	$\langle \Delta R_{\text{Truth}}^{\text{SIFT}} \rangle$	$\sigma_{\Delta R}^{\text{SIFT}}$	$\langle \Delta R_{\text{Truth}}^{\text{DELPHES}} \rangle$	$\sigma_{\Delta R}^{\text{DELPHES}}$
100	-0.059	0.054	+0.034	0.052	0.038	0.052	0.043	0.049
200	-0.052	0.049	+0.013	0.040	0.030	0.050	0.030	0.050
400	-0.045	0.045	+0.004	0.033	0.025	0.047	0.022	0.047
800	-0.040	0.042	-0.001	0.031	0.022	0.045	0.018	0.046
1600	-0.035	0.039	-0.003	0.029	0.018	0.040	0.015	0.042
3200	-0.031	0.035	-0.004	0.025	0.015	0.032	0.011	0.033

effects, whereas the angular response is improved primarily by the elimination of initial-state radiation. The most distinctive difference between the SIFT and large-radius DELPHES jets at this level is that the former is bounded from above by the partonic p_T , while the latter commonly exceeds it. The observed momentum excess is attributable to the capture of radiation from the underlying event. However, SIFT's filtering stage is apparently more adept at rejecting this contaminant, producing a reflection in the tail orientation that is reminiscent of various approaches to grooming.

Proceeding, we turn to attention the reconstruction of W -boson and top quark mass resonances, as visualized in Fig. 13 at each of the four central simulated p_T ranges. M_W

and M_t are respectively recovered from di- and trijet samples, by summing and squaring residual four-vector components after filtering. A second W reconstruction is obtained from decays of a t by optimizing the combinatoric selection of two prongs from the ($N = 3$) clustering flow. An excess near $M_W \simeq 80$ GeV is apparent for $p_T \geq 200$ GeV, although the top quark remains unresolved by the leading SIFT jet at low boost, since the associated bottom is likely to be separately isolated. The top-quark bump is clearly visible for $p_T \geq 400$ GeV, though its centroid falls somewhat to the left of $M_t \simeq 175$ GeV. The plotted distributions narrow at higher boost, and substantially sharper peaks are observed for $p_T \geq 800$ GeV. The systematic underestimation of mass is consistent with effects

FIG. 13. Distribution of W -boson and top quark masses for di- and tri-jet samples reconstructed with SIFT at various transverse boosts.




FIG. 14. Evolution of the measure δ_{AB}^N as a function of the number N of unmerged objects for QCD multijet production at various values of $\sqrt{\hat{s}}$. Partons have an angular separation ($\Delta R \geq 2.0$) and the p_T threshold is stepped in proportion to $\sqrt{\hat{s}}$. Initial-state radiation is suppressed, and analysis is at generator level. The orange band indicates the interval where all samples have merged to the point of their natural partonic count (black) but not beyond (white). The gray dashed line marks the isolation threshold at ($\delta_{AB} = 1$).

observed previously in the jet energy response, and it is similarly expected to be improvable with a suitable calibration.

IX. STRUCTURE TAGGING

This section describes applications of the SIFT algorithm related to structure and substructure tagging, including discrimination of events with varying partonic multiplicities, N subjettness axis finding, and identification of heavy resonances. We comparatively assess SIFT's performance on Monte Carlo collider data against standard approaches and quantify its discriminating power with the aid of a boosted decision tree (BDT).

Our first objective will be characterizing distinctive features in the evolution of the SIFT measure δ_{AB} for events with different numbers of hard prongs. We proceed by simulating pure QCD multijets representing LHC production of ($N = 2-5$) gluons and/or light first-generation quarks (u, d). Samples are generated at various partonic center-of-momentum energies, taking $\sqrt{\hat{s}} = (100, 200, 400, 800, 1600, 3200)$ GeV $\pm 20\%$. In order to ensure

that splittings are hard and wide (corresponding to a number of nonoverlapping large-radius jets with more or less commensurate scales), we require ($p_T \geq \sqrt{\hat{s}} \div 16$) and ($\Delta R \geq 2.0$). We also suppress initial-state radiation so that consistent partonic multiplicities can be achieved, and return to the use of generator-level (PYTHIA8) objects. Other selections and procedures are carried forward.

In contrast to the examples in Sec. VIII, the relevant showering products of these nonresonant systems are not expected to be captured within a single variable large-radius jet. Accordingly, we do not engage the isolation criterion from Sec. V for this application but instead apply exclusive clustering to the event as a whole with termination at ($N_{\text{exc}} = 1$). Filtering of soft-wide radiation is retained, but values of δ_{AB} are registered only for the merger of objects surviving to the final state. Relative to Fig. 4, candidate object pairings in the green and blue regions proceed to merge, while the softer member is discarded for those in the red region.

Figure 14 follows evolution of the SIFT measure as it progresses from 8 down to 1 remaining objects. Since we are considering entire events (as opposed to a hadronic




FIG. 15. Distribution of τ_2/τ_1 computed with SIFT axes for mono-, di-, and trijet samples at various transverse boosts.

event hemisphere recoiling off a neglected leptonic hemisphere), attention is focused here on the upper four values of $\sqrt{\hat{s}}$ to promote closer scale alignment with prior examples. Each of the simulated partonic multiplicities are tracked separately, represented by the geometric mean of δ_{AB}^N over all samples at level N in the clustering flow.

The relative change in the measure is larger when merging objects associated with distinct hard partons, suggesting that the jettness count is intrinsically imprinted on the clustering history. Specifically, a steepening in the log-slope of the measure evolution occurs when transitioning past the natural object count,⁸ i.e., from the black markers to the white markers. This supports the argument from Sec. VI that the most useful halting criterion can sometimes be none at all. In other words, it suggests that a determination of which objects should be considered resolved might best be made after observing how those objects would otherwise recombine.

The orange bands in Fig. 14 mark the range of δ_{AB}^N wherein all structures are fully reconstructed but not

⁸If the isolated objects have dissimilar p_T , then the discontinuity can be less severe, but the increased slope may extend to ($N = 1$).

overmerged, and the gray dashed line marks ($\delta_{AB} = 1$). Independently of the collision energy, white points tend to land above this line and black points below it, which helps to substantiate the isolation protocol from Sec. V. Bulk features of the evolution curves are substantially similar across the plotted examples and practically identical above the isolation cutoff, reflecting the scale-invariant design. However, δ_{AB}^N “starts” with a smaller value from large N for harder processes, and the orange band gap is expanded accordingly. This is because the tighter collimation [or smaller m/p_T , cf. Eq. (11)] associated with a large transverse boost induces smaller values of the measure when constituents are merged. Universality at termination is clarified by example, considering a true dijet system with balanced p_T , for which Eq. (15) indicates that ($\delta_{AB}^{N=1} \simeq \Delta R^2/2$). This is consistent with illustrated values around 20 for ($|\Delta\eta| \simeq 6$) and ($\Delta\phi = \pi$).

Our next objective will be to identify and test applications of SIFT for resolving substructure within a narrowly collimated beam of radiation. N subjettness represents one of the most prominent contemporary strategies for coping with loss of structure in boosted jets. In this prescription, one first clusters a large-radius jet, e.g., with ($R_0 \simeq 1.0$), which is engineered to contain all of the products of a decaying parton such as a boosted top or W boson. For




FIG. 16. Distribution of τ_3/τ_2 computed with SIFT axes for mono-, di-, and trijet samples at various transverse boosts.



various hypotheses of the subjet count ($N = 1, 2, 3, \dots$), a set of spatial axis directions are identified via a separate procedure, e.g., by reclustering all radiation gathered by the large-radius jet with an exclusive variant of the k_T or Cambridge-Aachen algorithms that forgoes beam isolation and forces explicit termination at N jets. One then computes a measure τ_N of compatibility with the hypothesis, which is proportional to a sum over minimal angular separation ΔR from any of the N axes, weighted by the transverse momentum p_T of each radiation component. Maximal discrimination of the subjet profile is achieved by taking ratios, e.g., τ_2/τ_1 or τ_3/τ_2 . This procedure will be our reference standard for benchmarking SIFT's substructure tagging performance.

The SIFT N -subjet tree automatically provides an ensemble of axis candidates at all relevant multiplicities that are intrinsically suitable for the computation of N subjeetiness. We test this claim using the previously described mono-, di-, and trijet event samples. The axis candidates are simply equal to the surviving objects at level N in the clustering flow. However, this process references only members of the leading isolated large-radius jet, rather than constituents of the event at large.

Figures 15 and 16 respectively exhibit distributions of τ_2/τ_1 and τ_3/τ_2 calculated in this manner at various

transverse boosts. The intuition that τ_3/τ_2 should be effective at separating W bosons from top quarks, whereas τ_2/τ_1 should be good for telling QCD monojets apart from W 's is readily validated. For comparison, Fig. 17 shows corresponding distributions of the same two quantities at the inner pair of p_T scales, as computed directly by DELPHES from the leading ($R_0 = 1.0$) soft-drop jet. Although there are qualitative differences between the two sets of distributions, their apparent power for substructure discrimination is more or less similar. This will be quantified subsequently with a BDT analysis.

Our final objective involves directly tagging substructure with sequential values of the SIFT measure. Distributions of δ_{AB} at the ($N = 1$), ($N = 2$), and ($N = 3$) clustering stages are plotted in Figs. 18, 19, and 20, respectively, for mono-, di-, and tri-jet samples at each of the four central simulated p_T ranges. Clear separation between the three tested object multiplicities is observed, with events bearing a greater count of partonic prongs tending to aggregate at larger values of the measure, especially after transitioning through their natural prong count. We observe that superior substructure discrimination is achieved by referencing the measure directly, rather than constructing ratios in the fashion

FIG. 17. Reference distributions of τ_2/τ_1 and τ_3/τ_2 computed by DELPHES at various transverse boosts.FIG. 18. Distribution of δ_{AB} at the $N = 1$ stage of clustering with SIFT for mono-, di-, and trijet samples at various transverse boosts.




FIG. 19. Distribution of δ_{AB} at the $N = 2$ stage of clustering with SIFT for mono-, di-, and trijet samples at various transverse boosts.

beneficial to N subjettness. This is connected to the fact that δ_{AB} is explicitly constructed as a ratio from the outset.

In order to concretely gauge relative performance of the described substructure taggers, we provide each set of simulated observables to a boosted decision tree for training and validation. BDTs are a kind of supervised machine learning that is useful for discreet (usually binary) classification in a high-dimensional space of numerical features. In contrast to “deep learning” approaches based around neural networks, where internal operations are shrouded behind a “black box” and the question of “what is learned” may be inscrutable, the mechanics of a BDT are entirely tractable and transparent. While neural networks excel at extracting hidden associations between “low level” features, e.g., raw image data at the pixel level, BDTs work best when seeded with “high-level” features curated for maximal information density.

At every stage of training, a BDT identifies which feature and what transition value optimally separates members of each class. This creates a branch point on a decision tree, and the procedure is iterated for samples following either fork. Classifications are continuous, typically on the range (0, 1), and are successively refined across a deep stack of shallow trees, each “boosted” (reweighted) to prioritize the correction of errors accumulated during prior stages.

Safeguards are available against overtraining on nonrepresentative features, and scoring is always validated on statistically independent samples. We use 50 trees with a maximal depth of five levels, a training fraction of $\frac{2}{3}$, a learning rate of $\eta = 0.5$, and L2 regularization with $\lambda = 0.1$ (but no L1 regularization). The BDT is implemented with MInOS [21], using XGBoost [30] on the backend.

The left-hand panel of Fig. 21 shows the distribution of classification scores for mono- and dijet event samples at $p_T = 1600$ GeV after training on values of the SIFT measure δ_{AB} associated with the final five stages of clustering. The two samples (plotted respectively in blue “background” and orange “signal”) exhibit clear separation, as would be expected from examination of the second element of Fig. 18. The underlying discretized sample data are represented with translucent histograms, and the interpolation into continuous distribution functions is shown with solid lines.

The right-hand panel of Fig. 21 shows the associated receiver operating characteristic (ROC) curve, which plots the true-positive rate versus the false-positive rate as a function of a sliding cutoff for the signal classification score. The area-under-curve (AUC) score, i.e., the fractional coverage of the shaded blue region, is a good proxy for overall discriminating power. A score of 0.5 indicates




FIG. 20. Distribution of δ_{AB} at the $N = 3$ stage of clustering with SIFT for mono-, di-, and trijet samples at various transverse boosts.

no separation, whereas classifiers approaching the score of 1.0 are progressively ideal.

The AUC (0.91) from the example in Fig. 21 is collected with related results in Table III. Separability of mono- and

dijet samples is quantified at each simulated range of p_T while making various feature sets available to the BDT. The first column uses the four DELPHES N -subjettiness ratios built from τ_1 to τ_5 . The next column references the same




FIG. 21. Left: Example distribution of BDT classification scores for the discrimination of mono- and dijet samples, respectively “background” and “signal,” at $p_T = 1600$ GeV. Training features include the δ_{AB}^N for ($N = 1-5$). Right: Associated receiver operating characteristic curve for true positives versus false positives.

TABLE III. Area under curve ROC scores for discrimination of resonances with hard 1- and 2-prong substructure using a BDT trained on various sets of event observables.

$p_T^{\text{GeV} \pm 5\%}$	$\tau_{\text{DELPHES}}^{N+1/N}$	$\tau_{\text{SIFT}}^{N+1/N}$	δ_{AB}^N	$\delta + \tau$
100	0.62	0.68	0.69	0.70
200	0.91	0.86	0.88	0.89
400	0.89	0.85	0.91	0.92
800	0.82	0.79	0.92	0.93
1600	0.77	0.74	0.91	0.92
3200	0.78	0.76	0.88	0.90

four ratios but as computed with objects and axes from the leading SIFT N -subjet tree. The third column provides the BDT with the final ($N = 1-5$) values of the SIFT measure δ_{AB}^N . The last column merges information from the prior two.

The two N -subjettiness computations perform similarly, but the fixed-radius DELPHES implementation shows an advantage of a few points in the majority of trials. The performance of N subjettiness degrades at large boost, losing around 10 points relative to its high scores by $p_T \simeq 1600$ GeV. The SIFT δ_{AB} measure outperforms N subjettiness in five of six trials, with an average advantage (over trials) of seven points. Its performance is very stable at larger boosts, where it has an advantage of at least ten points for $p_T \geq 800$ GeV. Combining the SIFT measure with N subjettiness generates a marginal advantage of about one point relative to δ_{AB} alone.

Table IV represents a similar comparison of discriminating power between resonances associated with hard 2- and 3-prong substructures. N subjettiness is less performant in this application, and the associated AUC scores drop by around six points. Performance of the SIFT measure degrades for soft events, but it maintains efficacy for events at intermediate scales and shows substantial improvement for $p_T \geq 800$ GeV, where its advantage over N subjettiness grows to around 20 points.

Table V extends the comparison to resonances with hard 3- and 1-prong substructures. The SIFT N -subjettiness computation is marginally preferred here over its fixed-radius counterpart. The δ_{AB} measure remains the best single

TABLE IV. Area under curve ROC scores for discrimination of resonances with hard 2- and 3-prong substructure using a BDT trained on various sets of event observables.

$p_T^{\text{GeV} \pm 5\%}$	$\tau_{\text{DELPHES}}^{N+1/N}$	$\tau_{\text{SIFT}}^{N+1/N}$	δ_{AB}^N	$\delta + \tau$
100	0.61	0.61	0.63	0.65
200	0.63	0.60	0.71	0.72
400	0.82	0.74	0.90	0.90
800	0.85	0.80	0.94	0.95
1600	0.77	0.77	0.97	0.97
3200	0.77	0.79	0.98	0.99

TABLE V. Area under curve ROC scores for discrimination of resonances with hard 3- and 1-prong substructure using a BDT trained on various sets of event observables.

$p_T^{\text{GeV} \pm 5\%}$	$\tau_{\text{DELPHES}}^{N+1/N}$	$\tau_{\text{SIFT}}^{N+1/N}$	δ_{AB}^N	$\delta + \tau$
100	0.70	0.75	0.77	0.77
200	0.86	0.87	0.90	0.90
400	0.93	0.91	0.95	0.96
800	0.91	0.89	0.96	0.96
1600	0.84	0.83	0.94	0.95
3200	0.76	0.78	0.91	0.92

discriminant by a significant margin, yielding an AUC at or above 0.90 for $p_T \geq 200$ GeV.

We conclude this section with a note on several additional procedural variations that were tested. Some manner of jet boundary enforcement (either via a fixed R_0 or the SIFT isolation criterion) is observed to be essential to the success of all described applications. Likewise, filtering of soft/wide radiation is vital to axis finding, computation of N subjettiness, and the reconstruction of mass resonances. Increasing the preclustering cone size from 0.01 radians to 0.1 substantially degrades the performance of N subjettiness, whereas discrimination with δ_{AB} is more resilient to this change.

X. COMPUTABILITY AND SAFETY

This section addresses theoretical considerations associated with computability of the SIFT observable δ_{AB} . Expressions are developed for various limits of interest. Infrared and collinear safety is confirmed, and deviations from recursive safety are calculated and assessed. It is suggested that SIFT's embedded filtering criterion may help to regulate anomalous behaviors in the latter context, improving on the Geneva algorithm.

Soft and collinear singularities drive the QCD matrix element governing the process of hadronic showering. In order to compare experimental results against theoretical predictions, it is typically necessary to perform all-order resummation over perturbative splittings. In the context of computing observables related to jet clustering, the calculation must first be organized according to an unambiguous parametric understanding of the priority with which objects are to be merged, i.e., a statement of how the applicable distance measure ranks pairings of objects that are subject to the relevant poles. Specifically, cases of interest include objects that are (i) mutually hard but collinear, (ii) hierarchically dissimilar in scale, and (iii) mutually soft but at wide angular separation. Pairs in the first two categories are likely to be physically related by QCD, but those in the third are not.

In order to facilitate considerations of this type, we outline here how the Eq. (15) measure behaves in relevant limits. The angular factor $\Delta \hat{R}_{AB}^2$ carries intuition for small

differences by construction [cf. Eq. (11)], and its dependence on aggregated mass has been further clarified in and around Eq. (12). We turn attention then to the energy-dependent factor e^{AB} , as expressed in Eq. (13), in two limits of interest. First, we take the case of hierarchically dissimilar transverse energies, expanding in the ratio ($\alpha \equiv E_T^A/E_T^B$) about 0,

$$2 \times e^{AB} \Rightarrow 2\alpha + \dots \quad (24)$$

Next, we expand for small deviations ($\zeta \equiv E_T^A/E_T^B - 1$) from matched transverse energies,

$$2 \times e^{AB} \Rightarrow 1 - \frac{\zeta^2}{2} + \dots \quad (25)$$

The SIFT algorithm is observed to be safe in the soft/infrared and collinear (IRC) radiation limits, because the object separation measure explicitly vanishes as ($\alpha \Rightarrow 0$) or ($\Delta R \Rightarrow 0$), up to terms proportional to the daughter mass squares [cf. Eq. (12)] in the latter case. This feature ensures that splittings at small angular separation or with hierarchically distinct transverse energies will be reunited during clustering at high priority.

With a clustering sequence strictly ranked by generated mass, JADE was plagued by an ordering ambiguity between the first and third categories described above, which presented problems for resummation. Geneva resolved the problem of mergers between uncorrelated mutually soft objects at wide separation in the same way that SIFT does, by diverging when neither entry in the denominator carries a large energy.

Yet, both SIFT and Geneva fall short of meeting the *recursive* IRC safety conditions described in Ref. [31] at the measure level. The challenge arises when a soft and collinear emission splits secondarily into a very collinear pair, as first observed in Ref. [32]. This scenario is visualized in Fig. 22, with hard object λ recoiling off a much softer emission κ (having $E_T^\kappa/E_T^\lambda \ll 1$) at a narrow




FIG. 22. Hard object λ emits a soft and collinear object κ at separation $\Delta\eta$, which experiences a secondary collinear splitting into a pair of objects with comparable hardness ($z \simeq 1/2$).

pseudorapidity separation ($\Delta\eta \ll 1$). Azimuthal offsets are neglected here for simplicity. The secondary radiation products are of comparable hardness for the situation of interest, carrying momentum fractions ($z \simeq 1/2$) and $(1-z)$ relative to their parent object κ .

It can be that the members of this secondary pair each successively combine with the hard primary object rather than first merging with each other. This ordering ambiguity implies that the value of the measure δ_{AB} after the final recombination of all three objects is likewise sensitive to the details of the secondary splitting. However, the mismatch is guaranteed to be no more than a factor of 2. Accordingly, this is a much milder violation than one associated with a divergence (as for JADE). While it does present difficulties for standard approaches to automated computation, it does not exclude computation.

We conclude this section by sketching the relevant calculation, translating results from Appendix F of Ref. [31] into the language of the current work. The secondary splitting is characterized by a parameter $\mu_\kappa^2 \equiv (m_\kappa/p_T^\kappa)^2$. We further apply the limit ($\mu_\kappa^2 \ll 1$), which implies ($E_T^\kappa \simeq p_T^\kappa$), and treat the radiation products of object κ as individually massless. The value of the measure for merging these objects is readily computed with Eq. (4), yielding ($\delta_{1-z}^z \simeq 2\mu_\kappa^2$). Note that the coefficient comes from the sum of squares in the measure denominator, in the limit of a balanced splitting. The merger of objects λ and κ (given prior recombination of the κ products) is best treated with Eq. (15), defining $\mu_\lambda^2 \equiv (m_\lambda/p_T^\lambda)^2$, and applying the limits in Eqs. (12) and (24), as follows:

$$\delta_\lambda^\kappa \simeq \left(\frac{E_T^\kappa}{E_T^\lambda} \right) \times [(\Delta\eta_\lambda^\kappa)^2 + \mu_\kappa^2 + \mu_\lambda^2]. \quad (26)$$

However, if the remnants of object κ instead combine in turn with object λ , then the final value of the measure (taking $z \geq 1/2$ without loss of generality) is instead ($\delta_\lambda^z \simeq z \times \delta_\lambda^\kappa$). In addition to that overall rescaling, the μ_κ^2 term from Eq. (26) is absent from the analogous summation in this context. If the κ splitting is hierarchically imbalanced (with $z \simeq 1$), then the secondary splittings are less resistant to merging first and the terminal measure value becomes insensitive to the merging order.

For balanced splittings, the physical showering history will be “correctly” rewound if ($\delta_{1-z}^z < \delta_\lambda^{1-z}$). But, there are no applicable kinematic restrictions enforcing that condition, and SIFT’s preference for associating objects at dissimilar momentum scales actually constitutes a bias in the other direction. On the other hand, the filtering criterion can help curb potential ambiguities in this regime. Specifically, the energy scale factor ($2 \times e_{1-z}^z \simeq 1$) associated with products of object κ will be subject here to the Eq. (25) limit. So, the “wrong” order of association is strongly correlated with cases where ($e_\lambda^{1-z} \ll 1$), since this is generally required in order to overcome the tendency for

strict collinearity ($\Delta\tilde{R}_{1-z}^z \ll \Delta\tilde{R}_\lambda^{1-z}$) in secondary splittings to commensurate momentum scales. In turn, this enhances the likelihood that the drop condition from Eq. (22) will veto any such merger. A full clarification of the SIFT filtering criterion's implications for recursive IRC safety is beyond our current scope, but is of interest for future work.

XI. CONCLUSIONS AND SUMMARY

We have introduced a new scale-invariant jet clustering algorithm named SIFT (scale-invariant filtered tree) that maintains the resolution of substructure for collimated decay products at large boosts. This construction unifies the isolation of variable-large-radius jets, recursive grooming of soft wide-angle radiation, and finding of subjet-axis candidates into a single procedure. The associated measure asymptotically recovers angular and kinematic behaviors of algorithms in the k_T family, by preferring early association of soft radiation with a resilient hard axis, while avoiding the specification of a fixed cone size. Integrated filtering and variable-radius isolation criteria resolve the halting problem common to radius-free algorithms and block assimilation of soft wide-angle radiation. Mutually hard structures are preserved to the end of clustering, automatically generating a tree of subjet axis candidates at all multiplicities N for each isolated final-state object. Excellent object identification and kinematic reconstruction are maintained without parameter tuning across more than a magnitude order of transverse momentum scales, and superior resolution is exhibited for highly boosted partonic systems. The measure history captures information that is useful for tagging massive resonances, and we have demonstrated with the aid of supervised machine learning that this observable has substantially more power for discriminating narrow 1-, 2-, and 3-prong event shapes than the benchmark technique using N subjetteness. These properties suggest that SIFT may prove to be a useful tool for the continuing study of jet substructure.

ACKNOWLEDGMENTS

The authors thank Bhaskar Dutta, Teruki Kamon, William Shepherd, Andrea Banfi, Rok Medves, Roman Kogler, Anna Albrecht, Anna Benecke, Kevin Pedro, Gregory Soyez, David Curtin, and Sander Huisman for useful discussions. The work of A. J. L. was supported in part by the UC Southern California Hub, with funding from the UC National Laboratories division of the University of California Office of the President. The work of D. R. was supported in part by DOE Grant No. DE-SC0010813. The work of J. W. W. was supported in part by the National Science Foundation under Grants No. NSF PHY-2112799 and No. NSF PHY-1748958. J. W. W. thanks the Mitchell Institute of Fundamental Physics and Astronomy and the Kavli Institute for Theoretical Physics for kind hospitality.

High-performance computing resources were provided by Sam Houston State University.

APPENDIX A: REVIEW OF COLLIDER COORDINATES

This appendix provides a brief pedagogical review of hadron collider coordinates. In this application, it is traditional to use a mapping from four-vector coordinates $P_\mu \equiv \{E, \vec{p}\}$ into coordinates $\{\eta, \phi, p_T, m\}$ that are well-behaved under Lorentz boosts along the longitudinal axis $\hat{z}$ of the beam. The pseudorapidity η (defined following) is a pure function of the zenith angle θ ,

$$\eta \equiv \frac{1}{2} \ln \left(\frac{|\vec{p}| + p_z}{|\vec{p}| - p_z} \right) \equiv -\ln \tan \left(\frac{\theta}{2} \right). \quad (\text{A1})$$

Forward (or backward) scattering correspond to η equals plus (or minus) infinity, while $\eta = 0$ represents entirely transverse scattering. The azimuthal angle ϕ measures orientation about the $\hat{z}$ axis. The transverse momentum p_T is the magnitude of the 3-vector momentum $\vec{p}$ projection perpendicular to the beam,

$$p_T \equiv \sqrt{p_x^2 + p_y^2}. \quad (\text{A2})$$

The final parcel of kinematic information is the Lorentz-invariant mass m , which is especially important for jets representing the composition of several lower level physical objects. The four-vector sum of individually massless objects may accumulate cancellation in the three-momentum that manifests as a non-negligible mass square in the invariant product,

$$p_\mu p^\mu \equiv E^2 - \vec{p} \cdot \vec{p} = m^2. \quad (\text{A3})$$

The quantity ΔR provides a radianlike measure of the relativistic “angular separation” between an object pair,

$$\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}. \quad (\text{A4})$$

The motivation for the definition in Eq. (A1) is that differences $\Delta\eta$ in pseudorapidity are “nearly” invariant under longitudinal boosts. To be precise, differences in the rapidity y (defined following) are a strict longitudinal invariant (as are the transverse coordinates p_T and ϕ), and y converges with η in the relativistic ($m \ll p_T$) limit,

$$\begin{aligned} y &\equiv \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \\ &= \ln \left(\frac{\sqrt{\cosh^2 \eta + \frac{m^2}{p_T^2}} + \sinh \eta}{\sqrt{1 + \frac{m^2}{p_T^2}}} \right). \end{aligned} \quad (\text{A5})$$

Figure 23 provides a visualization of angular dependence of the pseudorapidity, along with deviation from the




FIG. 23. The pseudorapidity η (bold, orange) is plotted as a function of the polar angle θ . For comparison, the longitudinal rapidity y (fine, blue) is also shown for various values of (m/p_T) , equal to $(\frac{1}{2}, 1, 2, 5, 10, 20)$ from top to bottom.

rapidity for various amounts of transverse boost. See Ref. [33] for an argument on the primacy of rapidity y .

APPENDIX B: SOFTWARE IMPLEMENTATIONS

This appendix describes two publicly available implementations of the SIFT algorithm. It also summarizes provided materials that facilitate reproduction of the key analyses in this manuscript. Finally, it outlines challenges to and plans for future integration with the *FastJet* contributions library.

The *AEACuS*, *RHADAManTHUS*, and *MInOS* packages respectively automate the processes of event analysis, visualization, and machine learning in a collider physics context. These tools are distributed and maintained at GitHub [21] by JWW, and inquiries are welcome. A quick-start tutorial (with a link to a video presentation) is additionally available at Ref. [34]. In brief, each of these programs is invoked from the command line and is interpreted with PERL 5.8+. Certain back-end features are implemented in Python 3, importing modules that include Matplotlib and XGBoost, as noted previously. All instructions regarding the computation of observables, application of event selection, generation of plots, and application of machine learning are specified in an associated card file, using a compact meta-language. Control cards used in the

preparation of this work are provided with its source package on the arXiv.

These programs are designed for easy integration with the standard *MadGraph/MadEvent*, *PYTHIA8*, and *DELPHES* chain, and cards are similarly included that document our approach to event production, generator-level selections, and the simulation of showering, hadronization, and detector effects. *AEACuS* autogenerates an extended-LHCO event record that bundles parton, hadron, and detector-level information (with weights) from the primary simulation chain for subsequent analysis. It further facilitates a variety of jet clustering and substructure applications, including an implementation of the SIFT algorithm. This usage is documented further in the example “cut” cards. The output is a space-delimited plain-text record of observables for each passing event, which serves in turn as an input to subsequent plotting and machine-learning operations.

The second existing public implementation of the SIFT algorithm is in the *Mathematica* notebook used here to produce jet clustering films and still frames. That notebook is likewise distributed on GitHub, bundled with the tools described prior. To run the notebook, simply place a suitable extended-LHCO event record into its working directory and “evaluate initialization cells.” User-adjustable parameters are documented in the notebook, for stipulating the clustering algorithm (members of the k_T family are also supported), a cone size (as applicable), any halting and filtering criteria, and whether ultra-soft ghost radiation should be included. An .mp4 film is typically output within a few minutes on a laptop computer, although running the notebook with ghost radiation enabled can be considerably more time consuming.

Additionally, a third implementation is planned that interfaces with *FastJet*, facilitating broader exploration of the SIFT algorithm within existing workflows. The *FastJet* implementation of Cacciari and Salam is famous for reducing the naïve $\mathcal{O}(N^3)$ runtime required for iterative pairwise jet clustering to $\mathcal{O}(N^2)$ or even to $\mathcal{O}(N \log N)$ in certain cases. It assumes a pairwise clustering measure $\delta_{ij} \equiv \min[f(p_T^i), f(p_T^j)] \times g(\Delta R_{ij})$ composed as the product of a kinematic function referencing the minimal value of $f(p_T^{i,j})$ over the pair times a “geometric” measure $g(\Delta R_{ij})$ that is typically a power of the pairwise angular separation. In contrast to the k_T family of measures, where factorization under the “FastJet Lemma” reduces neighbor finding to the ΔR plane, the search for SIFT neighbors is necessarily active in (at least) three dimensions. Moreover, the associated measure [cf. Eq. (15)] prioritizes *furthest* neighbors along one of these axes (u). Adaptations capable of confronting these unique challenges while maintaining “linearithmic” $\mathcal{O}(N \log N)$ scaling will be presented in a future work.

- [1] M. Cacciari, G. P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [2] J. Thaler and K. Van Tilburg, Identifying boosted objects with N-subjettiness, *J. High Energy Phys.* **03** (2011) 015.
- [3] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, Soft drop, *J. High Energy Phys.* **05** (2014) 146.
- [4] I. W. Stewart, F. J. Tackmann, J. Thaler, C. K. Vermilion, and T. F. Wilkason, XCone: N-jettiness as an exclusive cone jet algorithm, *J. High Energy Phys.* **11** (2015) 072.
- [5] J. Thaler and T. F. Wilkason, Resolving boosted jets with XCone, *J. High Energy Phys.* **12** (2015) 051.
- [6] I. W. Stewart, F. J. Tackmann, and W. J. Waalewijn, N-Jettiness: An Inclusive Event Shape to Veto Jets, *Phys. Rev. Lett.* **105**, 092002 (2010).
- [7] S. Bethke, Z. Kunszt, D. Soper, and W. Stirling, New jet cluster algorithms: Next-to-leading order QCD and hadronization corrections, *Nucl. Phys.* **B370**, 310 (1992).
- [8] M. Gouzevitch, A. Oliveira, J. Rojo, R. Rosenfeld, G. P. Salam, and V. Sanz, Scale-invariant resonance tagging in multijet events and new physics in Higgs pair production, *J. High Energy Phys.* **07** (2013) 148.
- [9] T. Lapsien, R. Kogler, and J. Haller, A new tagger for hadronically decaying heavy particles at the LHC, *Eur. Phys. J. C* **76**, 600 (2016).
- [10] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [11] M. Cacciari and G. P. Salam, Dispelling the N^3 myth for the k_t jet-finder, *Phys. Lett. B* **641**, 57 (2006).
- [12] S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber, Longitudinally invariant K_t clustering algorithms for hadron hadron collisions, *Nucl. Phys.* **B406**, 187 (1993).
- [13] S. D. Ellis and D. E. Soper, Successive combination jet algorithm for hadron collisions, *Phys. Rev. D* **48**, 3160 (1993).
- [14] Y. L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber, Better jet clustering algorithms, *J. High Energy Phys.* **08** (1997) 001.
- [15] M. Wobisch and T. Wengler, in *Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting)* (DESY, Hamburg, Germany, 1998), pp. 270–279, [arXiv: hep-ph/9907280](#).
- [16] S. Bethke *et al.* (JADE Collaboration), Experimental investigation of the energy dependence of the strong coupling strength, *Phys. Lett. B* **213**, 235 (1988).
- [17] W. Bartel *et al.* (JADE Collaboration), Experimental studies on multi-jet production in e^+e^- annihilation at PETRA energies, *Z. Phys. C* **33**, 23 (1986).
- [18] F. A. Dreyer, L. Necib, G. Soyez, and J. Thaler, Recursive soft drop, *J. High Energy Phys.* **06** (2018) 093.
- [19] J. M. Butterworth, A. R. Davison, M. Rubin, and G. P. Salam, Jet Substructure as a New Higgs Search Channel at the LHC, *Phys. Rev. Lett.* **100**, 242001 (2008).
- [20] J. W. Walker, Large hadron collider jet clustering visualizations (2022), <https://youtube.com/playlist?list=PLWgaPsMt19Rsu53Wh7r-rj8oF-rX-tDOt>.
- [21] J. W. Walker, AEACuS, RHADAManTHUS, and MinOS: Automated tools for collider event analysis, plotting, and machine learning (2023), <https://github.com/joelwwalker/AEACuS>.
- [22] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [23] T. Sjstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [24] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi (DELPHES 3 Collaboration), DELPHES 3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [25] M. Cacciari, G. P. Salam, and G. Soyez, The catchment area of jets, *J. High Energy Phys.* **04** (2008) 005.
- [26] D. Bertolini, P. Harris, M. Low, and N. Tran, Pileup per particle identification, *J. High Energy Phys.* **10** (2014) 059.
- [27] D. Krohn, J. Thaler, and L.-T. Wang, Jet trimming, *J. High Energy Phys.* **02** (2010) 084.
- [28] S. D. Ellis, C. K. Vermilion, and J. R. Walsh, Recombination algorithms and jet substructure: Pruning as a tool for heavy particle searches, *Phys. Rev. D* **81**, 094023 (2010).
- [29] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [30] T. Chen and C. Guestrin, in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York, NY, USA, 2016), KDD '16, pp. 785–794, ISBN 978-1-4503-4232-2.
- [31] A. Banfi, G. P. Salam, and G. Zanderighi, Principles of general final-state resummation and automated implementation, *J. High Energy Phys.* **03** (2005) 073.
- [32] S. Catani, B. R. Webber, Y. L. Dokshitzer, and F. Fiorani, Average multiplicities in two and three jet e^+e^- annihilation events, *Nucl. Phys.* **B383**, 419 (1992).
- [33] J. Gallicchio and Y.-T. Chien, Quit using pseudorapidity, transverse energy, and massless constituents, [arXiv:1802.05356](#).
- [34] J. W. Walker, Automated collider event selection, plotting, & machine learning with AEACuS, RHADAManTHUS, & MinOS, *Proc. Sci. CompTools2021* (2022) 071.